Center for High Performance Computing NEWS

THE UNIVERSITY OF UTAH SPRING 2011

Research

CHPC Supports Biomedical Research on Epigenetics

by Janet Ellingson

We know that our physical appearance is due in large part to the genes we inherit from our parents. We also know that alterations in a DNA sequence may result in developmental abnormalities and cancer. In addition to the DNA sequence of genes, the proper expression of the gene (whether and when a gene is turned off or on) also affects human development. Much of the field of epigenetics involves the study of the mechanisms that affect gene expression. With

Researchers of the Cairns Lab, Huntsman Cancer Institute

the aid of CHPC computing resources, Brad Cairns and his research group at the Huntsman Cancer Institute are investigating the role of chromatin structure – the packaging of DNA - in regulating gene expression.

Chromatin is the protein components that package regions of DNA into chromosomes. Stretched end-to-end, there is about 6 feet of DNA in each human cell, which must be packaged down to micron size by a set of protein molecules termed histones. These histones combine with the DNA to form bead-like nucleosomes that resemble 'beads on a string' when assembled along chromosomes, and they form the main component of chromatin. These nucleosome beads then coil into higher level structures that fit snuggly within the cell's nucleus. These nucleosome beads are also marked by chemical tags to help the gene expression machinery find needed genes. Exactly how nucleosomes are structured, positioned on genes, and chemically marked matters. Dr. Cairns observed that "certain chromatin structures silence genes, whereas others facilitate gene expression." The goals of his research lab are "to understand how the packaging by chromatin influences transcription and to characterize the protein machines that bind those chemical tags and facilitate the transitions between chromatin states."

Cells have "remodelers," an appropriate name for the biological machines that control the spacing and density of nucleosomes within the chromatin, allow access to the DNA and restructure the chromatin to create specialized regions. Genetic studies have shown that remodelers, composed primarily of proteins, have important roles in the regulation of gene expression, allowing for normal body and organ development and cell differentiation.

(cont. on page 2)

New Services: Consulting for and Hosting of Research Data Bases

by Janet Ellingson and Jody Smith

CHPC is pleased to announce the introduction of research database support as a new service to University faculty and students. This new service includes data analysis, database design and development, database administration, and support for research database issues and concerns.

Jody Smith, CHPC database administrator, is available to provide these services and support. Jody, who recently joined the CHPC staff, has extensive experience in data warehousing, relational database design, development and implementation. He has worked with large commercial companies in the past and is now looking forward to helping University researchers develop and manage their research databases.

CHPC now provides the following services:

Data Analysis: Consultation for data analysis, modeling, and mapping is available. Training and database design concepts documentation can be requested if necessary.

(cont. on page 2)

(Database cont.)

Database Design and Development: These services entail creating a database schematic and tables that map to and contain the data defined in the analysis phase, importing and possibly creating an importing process if necessary, and validating and testing the populated database.

Database Administration: CHPC provides training, maintenance (including backups), reorganizing and optimization, security, upgrades and deployment.

Other database services may include technical advice on performance and compatibility issues of existing databases, database conversion implications and considerations, as well as client and web interface products that allow easier database access, enhancements and modifications.

Standards and Tools: The standards and tools used for this support will be mainstream technologies which have been tested extensively and have been reliable and stable for today's vast amount of data and media storage, retrieval, and interpretation needs. These standards include UML data model structures, relational and object oriented methodologies and database management systems. The tools available are UNIX Servers, MS Access, MySQL, Postgres, UML Designer, XML, PHP, Python, and .Net.

An example of a CHPC-supported project is the design, data modeling and database development for a research project conducted by Professors Tom Kursar and Phyllis Coley of Biology. They and their students have been studying plant taxonomy (the defenses plants use against herbivores), specifically targeting the "Inga" species found in South America. The collected data includes plant chemicals, herbivore taxonomy, photos of plants and herbivores, and ants that visit nectarines. At this point, they have identified, defined, and gathered 85% of the data, created a data model and database schema, and populated 75% of the database.

If CHPC database services would help you manage your research data, please contact Jody Smith. He can be reached by phone (801-647-3042) or by email: jody.d.smith@utah.edu.

Jody Smith, CHPC's database administrator

(Biomedical research cont.)

Therefore, a remodeler's misregulation may result in human syndromes and cancers. Understanding the regulatory activity of remodelers, an area of increasing complexity, is an essential step toward a more thorough understanding of how to prevent genetic defects and disease.

The Cairns group is interested in whether chromatin structures in the germ cells of one generation (sperm) can influence the development of the embryo, and is currently focusing on the distinctive chromatin regulators in human sperm. During sperm creation, most of the nucleosomes found in chromatin are replaced with protamines. Given the importance of nucleosomes in the chromatin packaging, and therefore in the gene expression, the provocative question arises about the role these remaining nucleosomes have in embryonic development. The lab's experiments involve isolating specific DNA fragments that are attached to nucleosomes of interest by doing tens of millions of DNA sequencing reactions on each sample to identify the DNA fragment attached to the protein. CHPC computing resources become involved when the lab uses fairly sophisticated computer programs to match the DNA sequences to the genome – a process called "alignment." This work is done in collaboration with the Bioinformatics Shared Resource at the Huntsman Cancer Institute. Dr. Cairns states "we are very grateful for the presence and services of the CHPC, which is catalyzing our work and the work of others in genomics."

In order to observe the developmental consequences of the altered chromatin structure in gene transcription, the Cairns lab is using the germ cells (sperm and egg), stem cells, and embryos of the zebrafish, which has the same developmental packaging as human stem cells. By remodeling the chromatin structure in a zebrafish stem cell and observing transcription and the subsequent development of the embryo, the lab will examine whether chromatin changes and their effects on gene expression have been passed on to a new generation.

For more information about the research being done at the Cairns lab please visit the website: http://www.huntsmancancer.org/group/cairns/intro.jsp

NEWS

Construction on New Data Center Begins

by Steven Corbato

The University of Utah is developing a new off-campus data center that will significantly increase its data storage and computational server capabilities. The data center will be housed in an existing industrial building south of downtown Salt Lake City. In the late 1990s, the local Coca Cola bottling company moved its plant elsewhere and sold the building to Worldcom, a telecommunications company. Worldcom gutted the interior and reinforced the structure with the plans to establish an earthquake-proof data center. However, when the internet/telecom bubble burst a decade ago, bankruptcy struck Worldcom and the 74,000 square foot building fell into disuse. During the recent economic downturn, the University was able to purchase the building. Part of the building immediately was converted into a consolidated storage facility for legacy medical records (for example, patient charts and imaging films) that are now in the process of being digitized.

The University recently began the construction of the interior rooms and the cooling and energy supply systems for the data center. The design team developed a plan that will unify in one secure, stable, and efficient space a majority of the computing and storage needs of the University, including the University Health Care system and the Utah Education Network (UEN), which provides Internet, application, and curriculum services for all public (K-12) and state-supported higher education in Utah. In the future, additional public sector partners are expected to utilize this facility. The design team was led by VCBO Architecture of Salt Lake City with engineering expertise provided by SmithGroup and Alfa Tech. Oakland Construction serves as the general contractor for the building. The State's Division of Facilities Construction and Management and the University's Campus Design and Construction office are playing leadership roles in the design and construction process.

A primary design objective is environmental sustainability achieved through minimizing the energy and water consumption. Even though Utah has low industrial electric power rates, the electric bill for the data center would be prohibitive without significant energy conservation. The building initially will house almost 2,000 servers each generating a substantial demand for electricity and creating a tremendous amount of heat. The design team, working with the University's Office of Sustainability, plans to leverage our high, cold desert environment and employ ambient air and water cooling that will reduce energy consumption significantly. Except during the hottest days of the year, outside air will be used to cool the data center. During cold month, server-produced warm air will be mixed in to maintain the building at the optimum temperature.

University servers are now housed in various locations around campus. CHPC has its servers in the University Information Technology data center in Research Park, the INSCC building and the Student Services building. In spite of the multiple locations, there is still not enough space for CHPC to expand efficiently. In addition, each of the three locations has its unique cooling and electrical supply challenges. With one well-planned and well-constructed center, CHPC will improve its ability to efficiently and reliably deliver its services. The amount of electrical power available for CHPC research computing purposes will increase by three fold to 1.15 megawatts. CHPC will move its servers to the facility during the spring of 2012.

UIT officials at data center ground breaking ceremony on April 7, 2011

Along with the improvements to the building's interior, the University and UEN are developing a state-of-the-art optical network that will physically connect the data center to the University and other key research locations around the Salt Lake Valley, as well as to USU and BYU. The University and UEN are leveraging conduits and optical fibers that run along UTA light rail lines and UDOT highways. They have received significant financial support to develop this network from the National Science Foundation and the National Telecommunications and Information Administration.

For CHPC Users:

EMBER
Job Scheduling Policy
effective April 1, 2011

by Julia Harrison, CHPC Associate Director

Recently we made some configuration changes to the ember cluster. Most significantly, the nodes purchased by Professor Phil Smith's research funds (nicknamed "smithp" nodes) will be run under a separate reservation in the system and will be limited to a max wall clock time of 24 hours (instead of the current 48 hour limit). Users are allowed to run preemptable work on these nodes, but will need to specifically target them, and will need to limit wall time to the 24 hour maximum. General users do not need to be out of allocation to use these nodes, and these hours will not go against their ember allocation.

General jobs with allocation are run in the same configuration on the 53 general nodes as before, except we will increase the max wall time limit to 72 hours. Freecycle jobs are limited to the 53 general nodes, have a max wall time of 72 hours and continue to be preemptable.

Below is a summary of the scheduling policies in place today on the Ember cluster. Watch our web site for updated cluster scheduling policies for all of our production clusters. There are subtle differences between each cluster's policies, but the philosophy, approach and user tools remain consistent. As always, if you have any questions or concerns, please let us know by emailing to: issues@chpc.utah.edu.

- No node sharing.

- Allocations on the general side will be handled through the regular CHPC allocation committee. Allocations on owner nodes will be at the direction of node owners.

- Best effort to allocate nodes of same CPU speed.

- Max time limit for jobs will be as outlined in the QOS definitions (see table).

- Scheduling is set based on a current highest priority set for every job, excluding freecycle jobs which are scheduled in backfill (bestfit) mode.

- Fairshare boost in priority at user level. Minimal boost to help users who haven't been running recently. Our Fair share window is two weeks.

- Expansion Factor small boost in priority as queue time increases. Ratio between requested wall time versus eligible queue time.

- Reward for parallelism. Set at the global level.

- Max idle jobs in queue per user set to 5. This does not limit the number of jobs a user submits. Only top 5 will be eligible to run and accrue queue time priority.

- Majority of a job's priority will be set based on a quality of service definition or QOS (see table).

EMBER QOS

* available to general users

	Reservation	Priority	Description
* general	General_53	++	- allocatioan required - preemptor status - MAX runtime 72 hours
smithp	SmithP_187	++	- allocation required - preemptor status - MAX runtime 24 hours
bolton	Bolton_8	++	- allocation required - preemptor status - MAX runtime 14 days
kaplan	Kaplan_12	++	- allocation required - preemptor status - MAX runtime 14 days
long	General_53	++	- allocation required - special permision required - limited to 2 nodes at any time - MAX runtime 14 days
lmsmithp	LMSmithP_2	++	- allocation required - special access restrictions - MAX runtime 7 days
* freecycle (preemptee) (on general reserved nodes only)	General_53	flat or zero	- out of allocation required - preemptee status - jobs killed when preemptor wants to run - MAX runtime 72 hours - run in "backfill" only (not priority - jobs limited to 25% per user in freecycle (12 nodes) when there is competition
* sg (smithp-guest) (preemptee)	SmithP_197	flat or zero	- requires smithp-guest account for access - all charges will go against smithp-guest account - preemptee status but treated as allocated work (reservations respected) - MAX runtime 24 hours (same as smithp)
og (owner guest) (preemptee)	Bolton_8 Kaplan_12, etc. (determined by owner)	flat or zero	- requires owner-guest account for access - all charges will go against owner-guest account - preemptee status but treated as allocated work (reservations respected) -MAX runtime 72 hours

EMBER Standing Reservations

Standing Reservation	Access	Accounts	Node/ core count	Acceptance Criteria
* General_53	general	<pi>	53/636	qos=general, freecycle
LMSmithP_2	restricrted	smithp-lm	2/24	Account =smithp-lm
* SmithP_187	restricted	smithp-em smithp-guest	187/2244	qos=smith, sg (smithp-guest account)
Bolton_8	restricted	bolton-em	8/96	Account =bolton-em
Kaplan_12	restricted	kaplan-em	12/144	Account =kaplan-em

Ember cluster - photo by Sam Liston

A Sample of Research Using CHPC Resouces

DeTar, C., et al, "QCD thermodynamics with nonzero chemical potential at Nt=6 and effects from heavy quarks," (with the MILC collaboration), Physical Review D 81 (2010) 11,114504.

Pu, Z., and L. Zhang, 2010: Validation of AIRS temperature and moisture profiles over tropical oceans and their impact on numerical simulations of tropical cyclones," J. Geophys. Res. 15 (2010) D24114, doi:10.1029/2010JD014258.

Borodin, O., et al, " Molecular Dynamics Simulation and Pulsed-Field Gradient NMR Studies of FSI and TFSI-based Ionic Liquids," J. Phys. Chem. B 114 (2010) 6786-6798.

D. Bedrov, et al. "Influence of Polarization on Structural, Thermodynamic, and Dynamic Properties of Ionic Liquids Obtained from Molecular Dynamics Simulations," J. Phys.Chem. B 114 (2010) 4984-4997.

Freedman, H., L. P. Huynh, L. Le, T.E. Cheatam III, J. Tuszynski, T. N. Truong, "Explicitly-Solvated Ligand Contribution to Continuum Solvation Models for Binding Free Energies: Selectivity of Theophylline Binding to an RNA Aptamer," Journal of Physical Chemistry B 114 (2010) 2227-2237.

Zhdanov, M.S., Cuma, M., and Ueda, T., "3D electromagnetic holographic imaging in active monitoring of sea-bottom geoelectric structures," in Kasahara, J., V. Korneev, and M. S. Zhdanov, (eds), *Active Geophysical Monitoring*, Elsevier (2010) 317-342.

Facelli, J., "Chemical shift tensors: Theory and application to molecular structural problems," Progress in Nuclear Magnetic Resonance spectroscopy 58 (2011) 176 - 201.

SC11
Seattle, WA
Connecting communities through HPC

CHPC will have a booth at SuperComputing 2011 in November in Seattle. Designed by Sam Liston, our booth will showcase CHPC supported research. SC12 will be held in Salt Lake City at the Salt Palace. This will be a great opportunity to show off our area's computing and networking strengths. If you would like your research highlighted, contact Sam.

News

CHPC Hosts DUEL•ALITY

by Beth Miklavcic

On February 25-March 6, the University of Utah Center for High Performance Computing hosted a telematic research project entitled Duel•Ality, created and performed by CHPC Visualization Specialists Jimmy and Elizabeth Miklavcic. This original project investigated concepts of relationships and personal duality.

Duel•Ality explored the many sides of the interactions that the two characters Duel, played by Jimmy Miklavcic, and Ality, played by Elizabeth Miklavcic, encounter. This telematic work focused on the relationship between two people and the way this duo encounters and functions in this technological world. It also looked at the relationship of man to woman, of human to computer, of software to hardware, of live creativity to projected elements, as well as the relationship between creative elements and how the audience perceives them.

The script developed by Elizabeth and Jimmy Miklavcic included a series of monologues along with dialogues between Duel and Ality as they wrestled with each other and their expressions, using computers as extensions of themselves. This telematic play highlighted the human side of technology, revealing that everything and everyone has at least two sides.

The project laid down the foundations of investigation into real-time, telematic cinema where the performers manipulated and controlled the lighting, audio, video, computer graphics and animation during the live performance. The actors, as their respective characters, interacted with the environmental, visual and aural experience of the performance.

This method integrated the various technologies directly into the performance, bringing them into the visual forefront. Three cameras, a video mixer, an audio mixer, lighting control board and rear projection were all part of the project setting. The characters, Duel and Ality, while interacting with each other, also directly interacted with the technology.

Duel•Ality was streamed via the Access Grid and the Darwin QuickTime Streaming Server. Virtual venues include the University of Utah's ArtGrid venue, the Marriot Library Island in Second Life, Another Language Performing Arts Company's website (anotherlanguage.org) and Utah's University Information Technology's Wowza media server.

Congratulations! In 2010 Beth and Jimmy Miklavcic received the "Best Paper Award in the Arts Category" for their second edition of "InterPlay: Performing on a High Tech Wire" from the Utah Academy of Sciences, Arts and Letters.

What is CHPC?

The Center for High Performance Computing provides large-scale computing resources to University faculty and research staff to facilitate their research. CHPC is located in the INSCC building (just north of the Park administration building) and is responsible for the operation, maintenance and upgrade of their computing resources housed at data centers in INSCC, SSB and Komas.

The projects currently supported by CHPC come from a wide array of University disciplines that require large capacity computing resources, both for calculating the solutions of large-scale, two and three dimensional problems and for graphic visualization of the results.

If CHPC resources would be of use in your research, please go to our website www.chpc.utah.edu for more information.

CHPC Staff Directory

Administrative Staff	Title	Phone*	Email	Location
Julio Facelli	Director	585-3791	julio.facelli@utah.edu	410 INSCC
Julia D. Harrison	Associate Director	585-3791	julia.harrison@utah.edu	430 INSCC
Guy Adams	Assistant Director, Systems	554-0125	guy.adams@utah.edu	424 INSCC
Joe Breen	Assistant Director, Networking	550-9172	joe.breen@utah.edu	426 INSCC
DeeAnn Raynor	Administrative Officer	581-5253	dee.raynor@utah.edu	412 INSCC
Janet Ellingson	Admin. Program Coordinator & Newsletter Editor	585-3791	janet.ellingson@utah.edu	405 INSCC

Scientific Staff	Expertise	Phone*	Email	Location
Wim Cardoen	Scientific Applications	971-4184	wim.cardoen@utah.edu	420 INSCC
Martin Cuma	Scientific Applications	587-7770	martin.cuma@utah.edu	418 INSCC
Byron L. Davis	Statistics	585-5604	byron.davis@utah.edu	416 INSCC
Julio Facelli	Molecular Sciences	585-3791	julio.facelli@utah.edu	410 INSCC
Derek Huth	Scientific Support	N/A	N/A	405-8 INSCC
Sean Igo	Natural Language Processing	N/A	sean.igo@utah.edu	405-16 INCSS
Anita Orendt	Molecular Sciences	231-2762	anita.orendt@utah.edu	422 INSCC
Ron Price	Software Engineer & Grid Architect	560-2305	ron.price@utah.edu	405-4 INSCC
Jody Smith	Database Manager	647-3042	jody.d.smith@utah.edu	405-12 INSCC

Technical Support Staff	Group	Phone*	Email	Location
Irvin Allen	Systems	231-3194	irvin.allen@utah.edu	405-40 INSCC
Tom Ammon	Network	674-9273	tom.ammon@utah.edu	405-22 INSCC
Robert Bolton	Systems	528-8233	robert.bolton@utah.edu	405-24 INSCC
Wayne Bradford	Systems	243-8655	wayne.bradford@utah.edu	405-41 INSCC
Erik Brown	Systems	824-4996	erik.brown@utah.edu	405-29 INSCC
Brandon Day	User Services	N/A	N/A	405-20 INSCC
Steve Harper	Systems	541-3514	s.harper@utah.edu	405-31 INSCC
Brian Haymore	Systems	558-1150	brian.haymore@utah.edu	428 INSCC
Eric Hughes	Systems	879-8449	eric.hughes@utah.edu	405-18 INSCC
Samuel T. Liston	Systems, Multimedia	232-6932	sam.liston@utah.edu	405-39 INSCC
Kyle McGuire	User Services	N/A	N/A	405-14 INSCC
Randy Madsen	Systems	N/A	randy.madsen @utah.edu	5700 HSEB
Jimmy Miklavcic	Multimedia	585-9335	jimmy.miklavcic@utah.edu	296 INSCC
Beth Miklavcic	Multimedia	585-1066	beth.miklavcic@utah.edu	111 INSCC
Michael Palmer	Systems	435-720-3261	michaelj.palmer@utah.edu	405-28 INSCC
David Richardson	Network	550-3788	david.richardson@utah.edu	405-38 INSCC
Walter Scott	User Services	309-0763	walter.scott@utah.edu	405-13 INSCC
Steve Smith	Systems	581-7552	steve.smith@utah.edu	405-25 INSCC
Josh Spolsdoff	Systems	N/A	N/A	405-19 INSCC
Neal Todd	Systems	201-1761	neal.todd@utah.edu	405-30 INSCC
Alan Wisniewski	Network	580-5835	alan.wisniewski@utah.edu	405-21 INSCC

*All phone numbers are preceded by area code 801 unless otherwise noted.

The University of Utah seeks to provide equal access to its programs, services, and activities to people with disabilities. Reasonable prior notice is needed to arrange accommodations.

THE UNIVERSITY OF UTAH

Center for High Performance Computing
155 South 1452 East, RM #405
SALT LAKE CITY, UT 84112-0190

Welcome to CHPC News!

If you would like to be added to our mailing list, please fill out this form and return it to:

Janet Ellingson
THE UNIVERSITY OF UTAH
Center For High Performance Computing
155 S 1452 E ROOM 405
SALT LAKE CITY, UT 84112-0190
FAX: (801)585-5366

(room 405 of the INSCC Building)

Name:
Phone:

Department or Affiliation:
Email:

Address:
(UofU campus or U.S. Mail)

Thank you for using our Systems!

Please help us to continue to provide you with access to cutting edge equipment.

ACKNOWLEDGEMENTS

If you use CHPC computer time or staff resources, we request that you acknowledge this in technical reports, publications, and dissertations. Here is an example of what we ask you to include in your acknowledgements:

"*A grant of computer time from the Center for High Performance Computing is gratefully acknowledged.*"

Please submit copies or citations of dissertations, reports, pre-prints, and reprints in which the CHPC is acknowledged to: Center for High Performance Computing, 155 South 1452 East, Rm #405, University of Utah, Salt Lake City, Utah 84112-0190